Algebraic and Differential Topology in Machine Learning and AI: Obstructions, Invariants, and Counts

Anonymous

Higher Categorical Semantics · Version 1.0, 19 November 2026

Abstract. A neural network is a continuous map, usually piecewise linear, and the classical theorems of algebraic and differential topology are statements about what continuous maps can and cannot do. We organize their consequences for machine learning under three headings, pairing each qualitative statement with a computed one. *Obstructions*: by the Borsuk–Ulam theorem an autoencoder for data on an n-sphere with an n-dimensional bottleneck has worst-case reconstruction error at least the radius, and this is attained; a network output that names a rotation by a unit quaternion jumps by 2 somewhere on every non-contractible loop of rotations, because $\pi_1(\mathrm{SO}(3)) = \mathbb{Z}/2$, while a six-dimensional code does not; every tangent feature field on a sphere-like mesh has singularities of total index 2; and a closed piecewise linear curve with N pieces winds fewer than $N/2$ times, while composition multiplies degrees, so that depth L reaches degree 2^L with a fixed number of pieces per layer; and a network no wider than its input, with an injective activation, cannot enclose a bounded class. *Invariants*: diffusion with the Hodge Laplacian of a complex converges to a space whose dimension is a Betti number, at the rate of the spectral gap, and with real coefficients it cannot tell the projective plane from a disc; the Vietoris–Rips complex of twelve points on a circle is a circle up to a third of a turn, then three 2-spheres, then a 5-sphere; and persistence diagrams move by at most the size of a perturbation, which the count of their points does not. *Counts*: a classifier of the n-cube with small decision regions has points where $n+1$ of them meet (13 to 18 triple points for random prototype classifiers of the square); the fixed points of an attractor network on a circle or a torus have total index $1-d$ and 0; the degree of a normalized output changes only when the output passes through zero. A scalar odd network on an origin-centred circle has $2, 6, 10, \ldots$ zeros if all of them are simple. All numerical statements are recomputed by a companion script with 25 checks. No new theorem is claimed.

1 Introduction

The companion volume [ADTHC] retells two monographs of Pontryagin: one on the homology of polyhedra, ending with the fixed point theorem of Lefschetz and Hopf, and one on the classification of maps of spheres by framed manifolds. This paper collects what those theories say about learning systems. The connection is direct. A feed-forward network is a continuous map; with rectifier activations it is piecewise linear, a simplicial map after subdivision [ABMM18]. Data are assumed to lie near a manifold or a polyhedron [Car09, EdHa10]. Latent codes, rotation heads, feature fields on meshes and recurrent attractors are maps to and from spaces whose topology is known. Theorems about continuous maps then become statements about architectures, and they have a particular character: they hold for every choice of weights.

We sort the consequences by the kind of statement.

Obstructions say that something cannot be done continuously: a manifold cannot be coded in too few dimensions, a rotation cannot be named by four numbers, a sphere cannot be combed, a narrow network cannot enclose a region [ZBLYL19, CWKW19, FdH$^+$18, Joh19].

Invariants are what a functor computes. Simplicial and Hodge-theoretic networks [EDS20, BFW$^+$21, Lim20, PBB$^+$24] compute homology with real coefficients; persistent homology [CEH07, NSW08] computes it across scales, where the complexes themselves have topology of their own [AdAd17].

Counts are forced numbers: of points where classes meet, of fixed points of a recurrent map, of preimages of a value.

What is and is not claimed. Every statement carries a status tag: classical, literature, reformulation, observation or programme. The theorems are classical, and several of the applications are in the literature cited; where a formulation or a constant seems not to be standard we say "observation". Tags such as [04] name checks in the companion script `verify-topology-learning.py` (Appendix A), which uses only the Python standard library, exact arithmetic for homology, and a fixed seed.

The checks are instances and not proofs. The models are small and exact; no trained network appears. The references were assembled from memory, except [ZBLYL19, AdAd17, PBB+24], which were checked online on 19 November 2026.

2 Setting

A ***network*** is a continuous map $f_w\colon \mathbb{R}^a \to \mathbb{R}^b$ depending continuously on weights w. Data lie on a compact subset $M \subset \mathbb{R}^a$. A ***code*** for M in dimension m is a continuous *encoder* $e\colon M \to \mathbb{R}^m$; a ***decoder*** is any map $d\colon \mathbb{R}^m \to \mathbb{R}^a$, and the ***worst-case error*** of the pair is $\sup_{x\in M} |d(e(x)) - x|$. A ***representation*** of a space Y in $\mathbb{R}^m$ is a continuous map $g\colon \mathbb{R}^m \supseteq D \to Y$ onto Y; it is ***continuous*** in the sense of [ZBLYL19] if it has a continuous right inverse $Y \to D$, so that a target in Y can be turned into a continuous regression target.

The tools are the functors of [ADTHC]: homology H_*, which sends a continuous map to a homomorphism and a homotopy to an equality; the degree of a map between spheres; the fundamental group; and the Lefschetz number $L(f) = \sum_k (-1)^k \operatorname{tr}(f_* \mid H_k)$. An obstruction is a diagram that such a functor forbids. The simplest is the one that proves Brouwer's theorem [ADTHC, Thm. 4.2]:

$$S^{n-1} \hookrightarrow D^n \xrightarrow{r} S^{n-1} \quad \rightsquigarrow \quad \mathbb{Z} \to 0 \to \mathbb{Z}, \qquad \text{id} \text{ on both}$$

A retraction of the ball onto its boundary would factor the identity of $\mathbb{Z}$ through 0.

3 Obstructions

3.1 Bottlenecks

Theorem 3.1 (Bottleneck bound). [reformulation; cf. Borsuk–Ulam [Mat03]] *Let M be the unit sphere $S^n \subset \mathbb{R}^{n+1}$, $e\colon S^n \to \mathbb{R}^n$ continuous and $d\colon \mathbb{R}^n \to \mathbb{R}^{n+1}$ arbitrary. Then*

$$\sup_{x\in S^n} |d(e(x)) - x| \;\geq\; 1,$$

and equality holds for e the projection to the first n coordinates and d the inclusion.

Proof. By the Borsuk–Ulam theorem there is x with $e(x) = e(-x)$. Let y be the common reconstruction. Then $|y - x| + |y + x| \geq |2x| = 2$, so one of the two errors is at least 1. For the projection the error is $|x_{n+1}| \leq 1$. □

For $n = 1$ the collision was located for 200 random trigonometric encoders, and for $n = 2$ for 40 random cubic encoders, to rounding (Figure 1). [O1, O2] For each of 100 random encoders of the circle, some fibre of the code, taken to one percent of its range, is contained in no disc of radius less than 1.0000: no decoder can do better than error 1 there. [O3] The theorem is about the worst case. The mean error can be made small, by a code that is injective except near one point; the bound says that the failure cannot be removed, only moved and made thin. For a general closed n-manifold the qualitative statement survives: by Brouwer's invariance of domain, the companion of the invariance of dimension [ADTHC, Thm. 4.3], it has no continuous injective code in $\mathbb{R}^n$. The constant is special to round spheres. This is the topological reason for latent spaces that are themselves spheres, tori or groups [DFdC+18, FdH+18].

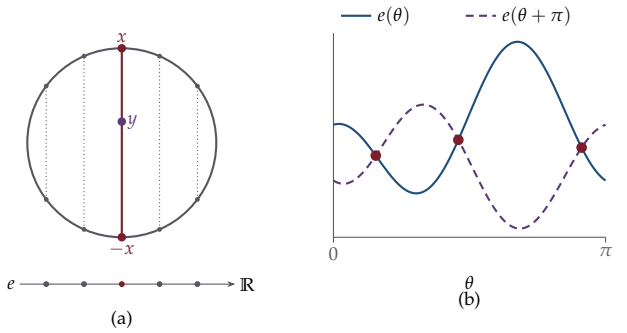


Figure 1: The bottleneck bound for the circle. (a) The fibres of the projection code; the fibre through the poles is an antipodal pair, and whatever point y the decoder returns for it is at distance at least 1 from one of the two. (b) A random trigonometric encoder e and its antipodal translate on a half-turn: since $e(\theta) - e(\theta + \pi)$ changes sign, the graphs cross (3 times here), and at a crossing the code identifies x and $-x$.

3.2 Rotations

Let $\rho\colon S^3 \to \mathrm{SO}(3)$ send a unit quaternion q to the rotation $v \mapsto qvq^{-1}$. It is onto, and $\rho(q) = \rho(-q)$.

Theorem 3.2 (Naming a rotation). [literature; [ZBLYL19]; [ADTHC], Thm. 10.1] (a) *ρ has no continuous right inverse: for every choice of a quaternion $s(R)$ with $\rho(s(R)) = R$ there is, on every non-contractible loop of rotations, a point where s jumps by $|q - (-q)| = 2$.* (b) *More generally a family of rotations $X \to \mathrm{SO}(3)$ has a continuous quaternion target if and only if it induces the zero map $\pi_1(X) \to \pi_1(\mathrm{SO}(3)) = \mathbb{Z}/2$.* (c) *There is no continuous representation of $\mathrm{SO}(3)$ in $\mathbb{R}^m$ for $m \leq 4$.* (d) *Two columns of the rotation matrix, with Gram–Schmidt orthogonalization as decoder, give a continuous representation in $\mathbb{R}^6$.*

Proof. (a), (b): ρ is the universal covering, and a map lifts through a covering exactly when it kills the fundamental group; a full turn about an axis lifts to a path from q to $-q$. (c) is proved in [ZBLYL19] from the fact that $\mathrm{SO}(3)$, which is real projective 3-space, does not embed in $\mathbb{R}^4$. (d) The two columns determine the third

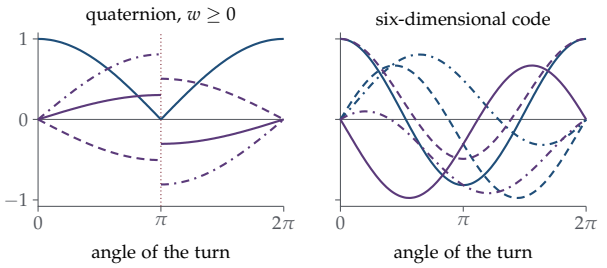


Figure 2: Two codes for the rotations about a fixed axis, through one full turn. Left: the four components of the unit quaternion with non-negative real part. At the half-turn the real part vanishes and the three others change sign, a jump of length 2.0000. Right: the six components of the first two columns of the rotation matrix, which return to their initial values without a jump.

by the cross product. □

Along a full turn in 2000 steps (Figure 2) the rotation matrix moves by at most 0.0044 per step and the six-dimensional code by 0.0040, while the quaternion with non-negative real part jumps by 2.0000 and the Euler angles of a turn about the second axis by 3.1416. [O4] For 200 random rotations the Gram–Schmidt decoder recovers the matrix to 10^{-12}, and it moves by at most 1.40 times a perturbation of the code. [O5] Statement (b) is the usable form: a data set of poses that contains a loop equivalent to a full turn has no continuous quaternion labelling, whatever the convention; a data set confined to a ball of rotations has one. A loss that identifies q with $-q$ repairs the target but not the output layer, which still has to choose.

3.3 Combing a mesh

A convolution on a surface compares feature vectors at neighbouring points, which requires a frame at each point in which to express them. A frame field is a framing in the sense of [ADTHC, Ch. 8].

Theorem 3.3 (Hairy ball, on meshes). [classical; Poincaré–Hopf; [Ban67]] *The zeros of a tangent vector field on a closed surface, counted with index, number χ. On a triangulated surface with a generic function h on the vertices, define the index of a vertex as $1 - \frac{1}{2}\#\{$sign changes of $h(w) - h(v)$ round the link$\}$; the indices add up to χ.*

The rule is drawn in Figure 3. The sums are $2, 2, 0, 1, 0$ for a tetrahedron, an icosahedron, the seven-vertex torus, the six-vertex projective plane and a Klein bottle, for 100 random functions each. [O6] So a global frame field exists on a closed orientable surface only when $\chi = 0$, and on a sphere every choice of frames is singular with total index 2. This is the obstruction that gauge-equivariant convolution is designed around [CWKW19]: since no frame can be chosen, the layers must be equivariant for changes of frame.

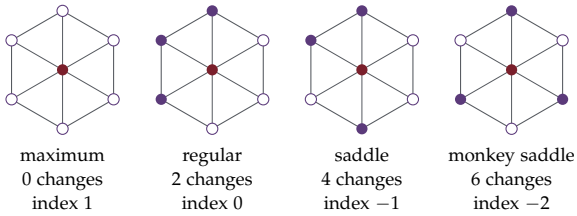


Figure 3: The index of a vertex for a function on a mesh. Filled neighbours are higher than the centre, open ones lower; the index is one minus half the number of changes round the link. A minimum has the same link as a maximum with the marks exchanged, and index 1.

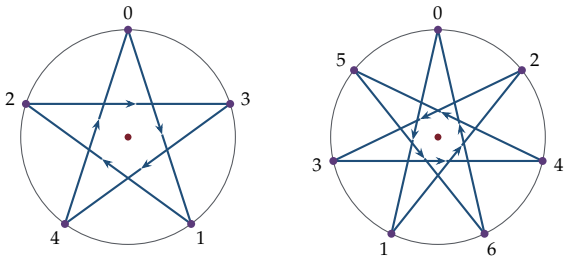
In higher dimensions the same question has a sharp answer. A sphere S^n with n even has $\chi = 2$ and not even one tangent field without zeros; and S^n has a global frame exactly for $n = 1, 3, 7$, a theorem of the same family as the Hopf invariant one problem [ADTHC, Rem. 9.7]. On S^3 the frame is explicit: at a unit quaternion q the vectors qi, qj, qk are orthonormal and tangent. [O11] This is why a latent space or an output space that is a circle, a torus or the group of unit quaternions needs no gauge, and why constructions that push a fixed noise distribution through the group multiplication are available there [FdH+18], while on S^2 they are not.

3.4 Pieces and degree

Theorem 3.4 (Capacity for winding). [observation] (a) *A continuous piecewise linear function of one variable with k breakpoints is the sum of an affine function and k rectifier units.* (b) *A closed piecewise linear curve in the punctured plane with N segments has winding number of absolute value less than $N/2$.* (c) *The polygon through the images of N equally spaced points under $z \mapsto z^d$ winds d' times, where $d' \equiv d \pmod N$ and $|d'| < N/2$.*

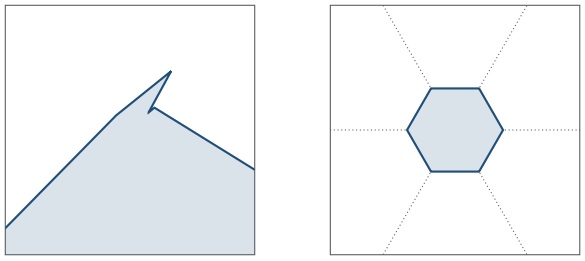
Proof. (a) To the leftmost affine piece add one unit $c_i \operatorname{relu}(x - x_i)$ per breakpoint x_i, with c_i the change of slope. (b) A segment missing the origin subtends an angle less than π. (c) Consecutive vertices differ by the angle $2\pi d/N$, taken in $(-\pi, \pi)$. □

Part (a) was confirmed on 50 random functions with six breakpoints. [O7] The polygons for $d = 3$ wind -2, 3, 3 times for $N = 5, 7, 12$ (Figure 4), and for $(d, N) = (5, 8)$ and $(7, 10)$ they wind -3 times. [O8] A network whose output curve has N linear pieces therefore cannot represent a map of the circle of degree d unless $N \geq 2|d| + 1$: too few pieces alias the degree exactly as too few samples alias a frequency. This is the simplicial approximation theorem [ADTHC, Thm. 3.2] with its hypothesis made visible: sampling a map at the vertices of a subdivision always gives a piecewise linear map, and that map is in the right homotopy class



$N = 5$: winding -2 $N = 7$: winding 3

Figure 4: The map $z \mapsto z^3$ of the circle sampled at N equally spaced points and joined by segments; the numbers give the order of the samples. With seven pieces each step is $\frac{3}{7}$ of a turn forwards and the polygon winds 3 times. With five, each step of $\frac{3}{5}$ of a turn is drawn as the shorter arc, $\frac{2}{5}$ of a turn backwards, and the polygon winds -2 times.



two units per layer three units, one layer

Figure 5: Decision regions on the square $[-3,3]^2$. Left: a network of depth three with two leaky rectifier units per layer; the positive region (shaded) is the image of a half-plane under a piecewise linear homeomorphism: it bends where units switch, and it has no holes and no end. Right: $1 - \sum_i \mathrm{relu}(a_i \cdot x)$ with three units; the dotted lines are where the units switch, and the positive region is a bounded hexagon.

only when the subdivision is fine enough for the star condition.

Proposition 3.5 (Depth multiplies degree). [reformulation; cf. [Tel16]] *Let F be the map of the plane that is linear on each of the eight cones over the half-sides of the square $[-1,1]^2$ and doubles the perimeter coordinate on its boundary. Then F^L maps the boundary to itself with degree 2^L and with $4 \cdot 2^L$ linear pieces. A network with one hidden layer of k rectifier units maps the boundary to a curve of at most $2k+4$ pieces, hence of degree less than $k+2$.*

Proof. Degree is multiplicative under composition [ADTHC, Prop. 5.4]. The breaking line of a unit meets the boundary of the square in at most two points, and the square has four corners; now apply Theorem 3.4*(b)*. □

The windings are $2, 4, \dots, 64$ and the pieces $8, 16, \dots, 256$ for $L = 1, \dots, 6$. [O9] Each layer is a fixed continuous piecewise linear map, hence a rectifier network of fixed size [ABMM18]. So a network of depth L and bounded width reaches degree 2^L, with a number of pieces within a factor 2 of the least that Theorem 3.4 allows, while one hidden layer needs more than $2^L - 2$ units. This is the depth-separation phenomenon [Tel16] with the count of oscillations replaced by a homotopy invariant: degrees multiply under composition and pieces only add within a layer.

3.5 Width

Theorem 3.6 (Narrow networks). [literature; [Joh19, HaSe17]] *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a network whose hidden layers all have n units, with invertible weight matrices, a non-zero output layer, and an activation that is continuous, strictly increasing and onto $\mathbb{R}$, such as the leaky rectifier. Then each decision region $\{f > 0\}$, $\{f < 0\}$ is homeomorphic to $\mathbb{R}^n$ and unbounded. In particular f cannot separate a bounded class from a class that surrounds it. With one hidden layer of $n+1$ rectifier units a bounded decision region exists.*

Proof. Each hidden layer is a homeomorphism of $\mathbb{R}^n$, so their composite h is one, and $\{f > 0\} = h^{-1}(H)$ for an open half-space H. If it were bounded its closure would be compact, and so would the image of the closure under h, which contains H. For the last statement let $a_0, \dots, a_n$ be the vertices of a regular simplex centred at the origin and $f(x) = 1 - \sum_i \mathrm{relu}(a_i \cdot x)$. Every $x \neq 0$ has $a_i \cdot x > 0$ for some i, and by compactness $\sum_i \mathrm{relu}(a_i \cdot x) \geq c|x|$ with $c > 0$; so $\{f > 0\}$ is bounded, and it contains the origin. □

For 30 random networks of depth four on the plane, with two leaky rectifier units per layer, two points of the same region were joined inside it by the preimage of a segment, the preimages of rays left every bounded set, and both classes were present on circles of radius 10, 100 and 1000 times the distance of the decision boundary from the origin. The three-unit function above is positive at the origin and at most -2.46 on the circle of radius 4 (Figure 5). [O10] The argument uses only that the layers are homeomorphisms: topology can change only where a layer folds, as the rectifier does, or changes dimension [NZL20]. For the rectifier itself and for other activations that are limits of injective ones a related statement holds, with unbounded level sets [Joh19]; and width $n+1$ is also the threshold for approximating arbitrary continuous functions [HaSe17].

4 Invariants

4.1 Hodge Laplacians

Let K be a simplicial complex with boundary matrices ∂_k. The ***Hodge Laplacian*** on k-chains is $L_k = \partial_k^\mathsf{T} \partial_k + \partial_{k+1} \partial_{k+1}^\mathsf{T}$; for $k = 0$ it is the graph Laplacian. Networks on complexes pass messages along ∂_k and $\partial_{k+1}^\mathsf{T}$ [EDS20, BFW+21].

Theorem 4.1 (Hodge). [classical; [Eck44, Lim20]] (a) $\ker L_k \cong H_k(K;\mathbb{R})$. (b) *The iteration* $X \mapsto X - \tau L_k X$ *with* $0 < \tau \le 1/\lambda_{\max}$ *converges to the orthogonal projection of* X *on* $\ker L_k$*, the error contracting by* $1 - \tau\lambda^+_{\min}$ *per layer, where* $\lambda^+_{\min}$ *is the smallest non-zero eigenvalue.*

The kernels have dimensions $(1,0,1)$, $(1,2,1)$, $(1,0,0)$, $(1,1,0)$, $(1,1,0)$, $(1,0,0)$ for the sphere, torus, projective plane, Klein bottle, Möbius band and disc, which are the Betti numbers computed independently by Smith normal form. [V1] On the edges of the seven-vertex torus $\lambda^+_{\min} = 1.5858$ and $\lambda_{\max} = 7$; after 99 layers four random inputs have converged to 10^{-11}, and their limits span a space of dimension 2. [V3] Statement (b) is the oversmoothing of graph networks one dimension up, with a difference: the limit is not a constant but a harmonic form, and it carries exactly the homology.

Proposition 4.2 (Coefficients are a choice of what is visible). [observation] *The limit in Theorem 4.1*(b) *depends only on the real homology of* K*. The projective plane and the disc have the same harmonic spaces in every degree, although* H_1 *of the projective plane is* $\mathbb{Z}/2$*. Over the field* $\mathbb{F}_2$ *the Betti numbers are* $(1,1,1)$ *and* $(1,0,0)$*. Conversely the Klein bottle and the torus agree over* $\mathbb{F}_2$*, with* $(1,2,1)$*, and differ over* $\mathbb{R}$*.*

The Betti numbers quoted are recomputed by the script. [V2] Integral homology sees both distinctions at once and is computed by integer row and column operations [ADTHC, Ch. 2], which are not differentiable. An architecture built from real boundary matrices has made the choice of the first column of the table below.

	over $\mathbb{R}$	over $\mathbb{F}_2$	over $\mathbb{Z}$
disc	$1,0,0$	$1,0,0$	$\mathbb{Z},0,0$
projective plane	$1,0,0$	$1,1,1$	$\mathbb{Z},\mathbb{Z}/2,0$
torus	$1,2,1$	$1,2,1$	$\mathbb{Z},\mathbb{Z}^2,\mathbb{Z}$
Klein bottle	$1,1,0$	$1,2,1$	$\mathbb{Z},\mathbb{Z}\oplus\mathbb{Z}/2,0$

4.2 Complexes built from samples

Topological data analysis replaces a point cloud by the ***Vietoris–Rips complex*** at scale r, whose simplices are the sets of points of diameter at most r, and reads off homology as r varies [Car09, EdHa10]. For small r and dense samples this recovers the homology of the underlying manifold [NSW08]. For large r it does not, and what it produces instead is known for the circle.

Theorem 4.3 (Rips complexes of a circle). [literature; [Ada13, AdAd17]] *As the scale increases, the Vietoris–Rips complex of the circle has the homotopy type of* S^1*, then* S^3*, then* S^5*, and so on, the type* S^{2l+1} *occurring for scales between* $\frac{l}{2l+1}$ *and* $\frac{l+1}{2l+3}$ *of a turn, until it is contractible. For* n *equally spaced points and scale* k *steps the same holds with* k/n *in place of the scale, and at the transition* $k/n = \frac{l}{2l+1}$ *the complex is a wedge of* $n - 2k - 1$ *spheres of dimension* $2l$*.*

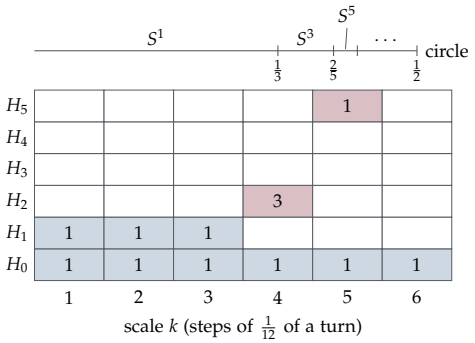


Figure 6: Betti numbers over $\mathbb{F}_2$ of the Vietoris–Rips complex of twelve equally spaced points on a circle, by scale. The class of the circle (blue) lasts up to $k = 3$. At $k = 4$, exactly a third of a turn, three classes appear in degree 2, and at $k = 5$ one in degree 5 (red); at $k = 6$ the complex is a full simplex. The line above shows the homotopy type of the complex of the whole circle at the corresponding scale [AdAd17].

For $n = 12$ the Betti numbers over $\mathbb{F}_2$ are $(1,1)$ for $k = 1,2,3$; $(1,0,3)$ for $k = 4$; and $(1,0,0,0,0,1)$ for $k = 5$ (Figure 6). For $n = 13$ and $k = 5$, which lies between $\frac{1}{3}$ and $\frac{2}{5}$, they are $(1,0,0,1)$. [V4, V5] For 16 random points on the unit circle whose largest gap is a chord of length 1.516, the complex has Betti numbers $(11,0)$ at $r = 0.15$, $(1,1)$ at $r = 1.55$ and $(1,0)$ at $r = 1.9$. [V6] The circle is seen from the largest gap up to about $\sqrt{3}$, the chord of a third of a turn. Above that scale the classes in dimensions 2, 3 and 5 are topology of the construction and not of the data. A learned model that consumes persistence diagrams should not treat high-dimensional classes at large scales as evidence of high-dimensional structure.

4.3 Stability

A ***persistence diagram*** records, for each homology class of the sublevel sets $\{f \le t\}$ of a function, the levels at which it is born and dies. The ***bottleneck distance*** between two diagrams is the least δ for which their points can be matched, to each other or to the diagonal, with no point moved by more than δ.

Theorem 4.4 (Stability). [classical; [CEH07]] *For two functions* f, g *on a finite complex, the bottleneck distance between their persistence diagrams is at most* $\max|f - g|$*.*

For 60 random functions on a cycle of 40 vertices, perturbed by between 0.02 and 0.3 in the maximum norm, the distance between the diagrams in degree 0 was never larger than the perturbation, and equal to it in the worst case, while the number of points of the diagram changed in 35 of the 60 cases. [V7] The diagram is a Lipschitz function of the data and a Betti number is not even continuous. This is the property that lets persistence enter a model as a feature or a loss, and it

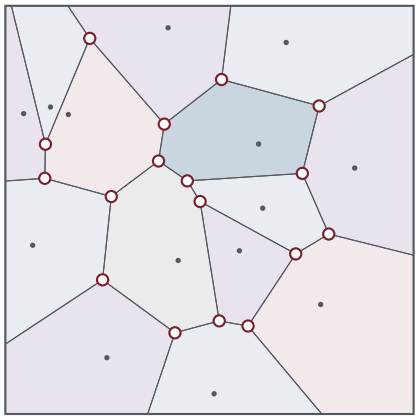


Figure 7: A nearest-prototype classifier of the square with 14 classes; the dots are the prototypes. No region reaches from one side to the opposite side, and there are 17 points (circled) where three classes meet. By Theorem 5.1 at least one is forced, however the regions are drawn.

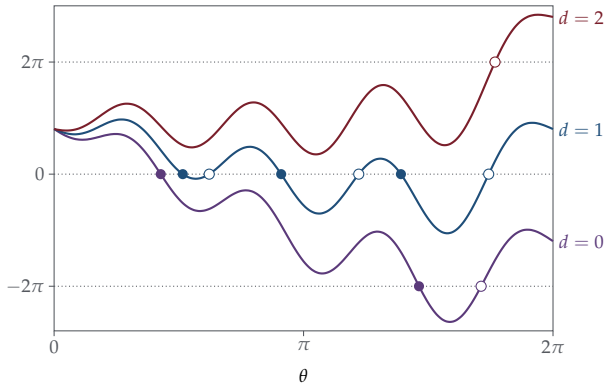


Figure 8: Fixed points of three maps $\theta \mapsto d\theta + \epsilon g(\theta)$ of the circle, with the same g. Plotted is the displacement $(d-1)\theta + \epsilon g(\theta)$; a fixed point is a crossing of a multiple of 2π (dotted). Downward crossings have index 1 (filled) and upward ones index -1 (open). The curve for d climbs by $d-1$ levels in one turn, so the indices add up to $1-d$: here 3, 6, 1 points with sums 1, 0, -1 for $d = 0, 1, 2$.

is the quantitative form of the functoriality of homology: a small change of f is an interleaving of the two filtrations.

5 Counts

5.1 Where regions meet

Theorem 5.1 (Junctions). [reformulation; cf. Lebesgue; [ADTHC], Thm. 4.3] *Let the cube $[0,1]^n$ be covered by finitely many closed sets, none of which meets two opposite faces. Then some point lies in at least $n+1$ of them.*

Apply this to the closures of the decision regions of a classifier, a region being a connected set of inputs with one label. If no region reaches across the cube, there is a point every neighbourhood of which meets $n+1$ regions. Small regions cannot be arranged so that at most n meet anywhere; the brick pattern, in which exactly $n+1$ meet, is optimal [ADTHC, Fig. 4.4].

For 60 random nearest-prototype classifiers of the square with 14 classes (one is drawn in Figure 7) there were between 13 and 18 triple points, and for 8 classifiers of the cube with 40 classes between 143 and 174 points where four classes meet. [N1, N2] The statement is a topological companion to results on the unavoidability of nearby misclassification [SHSFG19]: those use isoperimetric inequalities to estimate the measure of the points near a boundary, and this one says that the boundaries themselves must meet, $n+1$ at a time. The proof is Sperner's lemma, and the lemma comes with an algorithm: walk from cell to cell through faces carrying a fixed set of n labels. For a labelling induced by a smooth self-map of the triangle the walk reaches a complete cell, and with it an approximate fixed point, in 13, 25, 50, 97 steps on grids of side 16, 32, 64, 128, linear in the side although the number of cells is quadratic.

[N3] In high dimension the walk can be exponentially long, and finding its end is complete for the class PPAD [Pap94, DGP09]; this is the complexity of computing an equilibrium of a game, or a fixed point of an implicit layer [BKK19], without further structure.

5.2 Fixed points of attractors

Recurrent networks that hold a continuous variable do so on a manifold of fixed points: a ring for a heading, a torus for a position [BBS95, Zha96, GHP+22]. A trained or perturbed network is a map of that manifold close to the identity, or to another map of known degree.

Theorem 5.2 (Lefschetz counts). [classical; [ADTHC], Thm. 6.2, Prop. 6.3] (a) *A map of the circle of degree d with non-degenerate fixed points has fixed points of total index $1-d$; for $d=1$ their number is even, and for a map close to the identity half of them are stable.* (b) *The zeros of a vector field on a torus have total index 0, and on a sphere 2.* (c) *A map of a closed manifold homotopic to the identity has Lefschetz number χ.*

For circle maps $\theta \mapsto d\theta + \epsilon g(\theta)$ with random g and $\epsilon = 2$, the fixed points number $5, 7, 6, 9, 8$ for $d = -2, 0, 1, 2, 3$ and their indices add up to $3, 1, 0, -1, -2$ (Figure 8 shows three maps of the same kind). [N4] For three drift fields on a torus there are four zeros each, a source, a sink and two saddles. [N5] The consequences differ by manifold. On a ring or a torus, noise in the weights breaks the continuum of fixed points into isolated ones, stable and unstable in balance, or removes all of them and leaves a uniform drift, which $\chi = 0$ permits. On a sphere that second outcome is impossible: every perturbation of a spherical attractor keeps a fixed point, because $\chi = 2$.

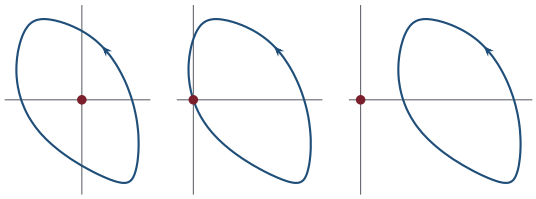


Figure 9: The output curve of a network with four hidden units, for three values of a bias b added to the first output; the dot is the origin. The normalized output has degree 1 while the curve encloses the origin and degree 0 after it has left; in between, at $b = 1.2930$, the curve passes through the origin and the normalized output is not defined.

5.3 Degree and the passage through zero

Let a network produce $u_w(\theta) \in \mathbb{R}^2$ for an input on a circle and output the direction $u_w/|u_w|$. As long as u_w does not vanish this is a map of the circle to itself with a degree.

Theorem 5.3. [classical; [ADTHC], Thm. 9.1] (a) *The degree is the signed number of preimages of any regular direction.* (b) *Along a path of weights the degree is constant unless $u_w(\theta) = 0$ for some θ and some weight on the path.*

For a network with four hidden units (Figure 9) the degree is 1, and the signed number of preimages of three directions is 1 each. Adding a bias b to the first output, the degree drops to 0 at $b = 1.2930$, where the smallest output norm over the circle is 10^{-4}. [N6] Training that must change the degree of the output, to fit a target of another class, has to pass through weights at which the normalized output is undefined and its gradient unbounded; and if the initial class is the right one, a barrier of the same kind protects it. By Hopf's theorem the degree is the only such invariant for maps to a sphere of the same dimension.

5.4 Zeros of odd networks

Proposition 5.4 (Odd networks). [classical; Borsuk–Ulam [Mat03]] *A network $\mathbb{R}^{n+1} \to \mathbb{R}^n$ without biases and with odd activations is an odd map, and it vanishes somewhere on every sphere centred at the origin. For $n = 1$, if its restriction to the circle is continuously differentiable and all its zeros are simple, their number is congruent to* 2 *modulo* 4.

Proof. Borsuk–Ulam gives an antipodal pair with equal values. Oddness makes that common value zero. For $n = 1$, start the half-turn at a point with nonzero output. The endpoint signs are opposite, and each simple zero changes the sign, so the half-turn contains an odd number of zeros. The other half-turn contains their antipodes, giving twice an odd number in all. □

A sign scan on the unit circle found 2 crossings for 193 of 200 random networks with six hyperbolic tangent units and 6 for the other 7. An output bias of 10 makes the last network strictly positive, since it exceeds the sum of absolute output weights. [N7] Oddness also applies to Theorem 3.1, whose collisions are the zeros of the odd map $e(x) - e(-x)$. An architecture made odd for the sake of a symmetry therefore pays for it: its scalar outputs cannot be of one sign on any sphere, and the decision boundary of a bias-free odd classifier meets every sphere round the origin.

6 Summary of results

Table 1 pairs the statements. The three headings correspond to three uses of a functor. An obstruction is a diagram that the functor forbids. An invariant is its value. A count is its value when that value is an integer that the geometry must realize as a number of points.

7 Limits

Worst case. Theorems 3.1 and 3.2 bound the supremum of an error. They say nothing about the average, and an obstruction confined to a set of small measure may be harmless in practice. What they exclude is its absence.

Idealization of data. The data are taken to be exactly a sphere, a rotation group or a circle. Real data are samples near such a set, and a network is evaluated on finitely many points, where every map is continuous. The theorems constrain the functions that interpolate, and through them the behaviour between samples.

Small exact models. The complexes have at most a few hundred simplices, the attractors are maps of a circle or a torus given by formulas, and the classifier is a nearest-prototype rule. No trained network was examined.

Real coefficients. Proposition 4.2 concerns the limit of diffusion. A simplicial network also sees the nonzero spectrum of its Laplacians, which can distinguish complexes with the same real homology; the statement is about harmonic features only.

Hypotheses on activations. Theorem 3.6 is proved for activations that are homeomorphisms of the line. For the rectifier the regions need not be homeomorphic to $\mathbb{R}^n$, and what survives is as cited.

Complexity. The PPAD-completeness quoted after Theorem 5.1 is for fixed-point problems given by circuits; the walk on a smooth example is short.

Novelty. The theorems are classical. Theorem 3.4, Proposition 3.5, Proposition 4.2 and the constant in Theorem 3.1 are tagged as observations or reformulations; they are elementary and may be in the literature in other words. The second monograph of the source, on framed manifolds and the first stable stems [ADTHC,

Topic	Qualitative statement	Quantitative statement	Checks
Bottleneck	no continuous injective code of S^n in $\mathbb{R}^n$	worst-case error ≥ 1, attained; fibre radius 1.0000 in 100 trials	O1–O3
Rotations	no continuous quaternion or Euler-angle target on a non-contractible loop; a six-dimensional one exists	jumps 2.0000 and 3.1416 against steps 0.004; decoder Lipschitz constant ≤ 1.40	O4, O5
Meshes	no global frame field unless $\chi = 0$; on spheres only for $n = 1, 3, 7$	index sums $2, 2, 0, 1, 0$ on five surfaces; frame on S^3 to 10^{-15}	O6, O11
Pieces	degree needs pieces	$N \geq 2\lvert d\rvert + 1$; windings $-2, 3, 3$ for $d = 3$, $N = 5, 7, 12$	O7, O8
Depth	composition multiplies degree	degree 2^L with $4 \cdot 2^L$ pieces at depth $L \leq 6$; one layer needs more than $2^L - 2$ units	O9
Width	n units per layer cannot enclose a region; $n + 1$ can	30 networks: both regions connected and unbounded; island value 1 inside, ≤ -2.46 outside	O10
Hodge diffusion	converges to real homology; blind to torsion	kernels equal Betti numbers on six complexes; gap 1.5858, 99 layers	V1–V3
Rips complexes	higher spheres at large scales are artefacts	$n = 12$ has S^1 for $1 \leq k \leq 3$, a wedge of three S^2 at $k = 4$, and S^5 at $k = 5$. Circle seen for $1.52 < r < \sqrt{3}$	V4–V6
Persistence	diagrams are stable, counts are not	distance $\leq$ perturbation in 60 trials, ratio up to 1.000; count changed in 35	V7
Decision regions	$n + 1$ regions meet	13–18 triple points; 143–174 quadruple points	N1, N2
Sperner walk	a complete cell is found by a path	$13, 25, 50, 97$ steps on sides $16, \ldots, 128$	N3
Attractors	index sum $1 - d$ on a circle, 0 on a torus, 2 on a sphere	$5, 7, 6, 9, 8$ fixed points with sums $3, 1, 0, -1, -2$; four zeros on the torus	N4, N5
Normalized outputs	degree changes only through zero	degree $1 \to 0$ at bias 1.2930, norm 10^{-4}	N6
Odd networks	a zero on every origin-centred sphere in $\mathbb{R}^{n+1}$ for n outputs	2 sign crossings in 193 and 6 in 7 of 200 networks. The last is positive with output bias 10	N7

Table 1: Qualitative statements with their quantitative counterparts. Every number is recomputed by the companion script.

Ch. 8–12], has no application here beyond the rotation groups, and we have not forced one.

A The companion script

`verify-topology-learning.py` runs in about ten seconds and prints one line per check. **O1–O3** antipodal collisions of codes on the circle and the sphere, and the radius of the worst fibre; **O4, O5** steps of four codes for rotations along a full turn, and the Gram–Schmidt decoder; **O6** index sums on five triangulated surfaces; **O7, O8** piecewise linear functions as rectifier networks, and winding numbers of polygons; **O9** windings and pieces of the iterated doubling map; **O10** decision regions of narrow networks, and a bounded region with three units; **O11** the quaternion frame on the 3-sphere; **V1–V3** Hodge Laplacians, coefficients, and convergence of diffusion; **V4–V6** Rips complexes of equally spaced and of random points on a circle; **V7** bottleneck distances between persistence diagrams of perturbed functions; **N1, N2** junctions of nearest-prototype classifiers in dimensions two and three; **N3** the length of the Sperner walk; **N4, N5** fixed points of circle maps and zeros of fields on a torus; **N6** the degree of a small network and the bias at which it changes; **N7** zeros of odd networks on a circle. Homology is computed exactly, by Smith normal form over the integers and Gaussian elimination over prime fields, with the library of the companion volume. Random inputs use fixed seeds. The figures were drawn by a separate plotting script from the same formulas and are embedded in the source as coordinates.

References

[ABMM18] R. Arora, A. Basu, P. Mianjy, A. Mukherjee, Understanding deep neural networks with rectified linear units, *Proc. ICLR* 2018.

[Ada13] M. Adamaszek, Clique complexes and graph powers, *Israel J. Math.* 196 (2013), 295–319.

[AdAd17] M. Adamaszek, H. Adams, The Vietoris–Rips complexes of a circle, *Pacific J. Math.* 290 (2017), 1–40. Checked.

[ADTHC] Anonymous, *Algebraic and Differential Topology in Higher Categories*, Higher Categorical Semantics, 2026. Companion volume.

[Ban67] T. Banchoff, Critical points and curvature for embedded polyhedra, *J. Differential Geom.* 1 (1967), 245–256.

[BBS95] R. Ben-Yishai, R. L. Bar-Or, H. Sompolinsky, Theory of orientation tuning in visual cortex, *Proc. Nat. Acad. Sci. U.S.A.* 92 (1995), 3844–3848.

[BFW+21] C. Bodnar, F. Frasca, Y. Wang, N. Otter, G. Montúfar, P. Liò, M. Bronstein, Weisfeiler and Lehman go topological: message passing simplicial networks, *Proc. ICML* 2021.

[BKK19] S. Bai, J. Z. Kolter, V. Koltun, Deep equilibrium models, *Adv. Neural Inf. Process. Syst.* 32 (2019).

[Car09] G. Carlsson, Topology and data, *Bull. Amer. Math. Soc.* 46 (2009), 255–308.

[CEH07] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Stability of persistence diagrams, *Discrete Comput. Geom.* 37 (2007), 103–120.

[CWKW19] T. Cohen, M. Weiler, B. Kicanaoglu, M. Welling, Gauge equivariant convolutional networks and the icosahedral CNN, *Proc. ICML* 2019.

[DFdC+18] T. R. Davidson, L. Falorsi, N. De Cao, T. Kipf, J. M. Tomczak, Hyperspherical variational auto-encoders, *Proc. UAI* 2018.

[DGP09] C. Daskalakis, P. W. Goldberg, C. H. Papadimitriou, The complexity of computing a Nash equilibrium, *SIAM J. Comput.* 39 (2009), 195–259.

[Eck44] B. Eckmann, Harmonische Funktionen und Randwertaufgaben in einem Komplex, *Comment. Math. Helv.* 17 (1944), 240–255.

[EdHa10] H. Edelsbrunner, J. Harer, *Computational Topology: An Introduction*, American Mathematical Society, 2010.

[EDS20] S. Ebli, M. Defferrard, G. Spreemann, Simplicial neural networks, NeurIPS Workshop on Topological Data Analysis and Beyond, 2020.

[FdH+18] L. Falorsi, P. de Haan, T. R. Davidson, N. De Cao, M. Weiler, P. Forré, T. S. Cohen, Explorations in homeomorphic variational auto-encoding, ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models, 2018.

[GHP+22] R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M.-B. Moser, E. I. Moser, Toroidal topology of population activity in grid cells, *Nature* 602 (2022), 123–128.

[HaSe17] B. Hanin, M. Sellke, Approximating continuous functions by ReLU nets of minimal width, arXiv:1710.11278, 2017.

[Joh19] J. Johnson, Deep, skinny neural networks are not universal approximators, *Proc. ICLR* 2019.

[Lim20] L.-H. Lim, Hodge Laplacians on graphs, *SIAM Rev.* 62 (2020), 685–715.

[Mat03] J. Matoušek, *Using the Borsuk–Ulam Theorem*, Springer, 2003.

[NSW08] P. Niyogi, S. Smale, S. Weinberger, Finding the homology of submanifolds with high confidence from random samples, *Discrete Comput. Geom.* 39 (2008), 419–441.

[NZL20] G. Naitzat, A. Zhitnikov, L.-H. Lim, Topology of deep neural networks, *J. Mach. Learn. Res.* 21 (2020), no. 184, 1–40.

[Pap94] C. H. Papadimitriou, On the complexity of the parity argument and other inefficient proofs of existence, *J. Comput. System Sci.* 48 (1994), 498–532.

[PBB+24] T. Papamarkou, T. Birdal, M. Bronstein, G. Carlsson, J. Curry, Y. Gao, M. Hajij, R. Kwitt, P. Liò, P. Di Lorenzo, V. Maroulas, N. Miolane, F. Nasrin, K. N. Ramamurthy, B. Rieck, S. Scardapane, M. T. Schaub, P. Veličković, B. Wang, Y. Wang, G.-W. Wei, G. Zamzmi, Position: topological deep learning is the new frontier for relational learning, *Proc. ICML* 2024, PMLR 235, 39529–39555. Checked.

[SHSFG19] A. Shafahi, W. R. Huang, C. Studer, S. Feizi, T. Goldstein, Are adversarial examples inevitable?, *Proc. ICLR* 2019.

[Tel16] M. Telgarsky, Benefits of depth in neural networks, *Proc. COLT* 2016, PMLR 49, 1517–1539.

[ZBLYL19] Y. Zhou, C. Barnes, J. Lu, J. Yang, H. Li, On the continuity of rotation representations in neural networks, *Proc. CVPR* 2019, 5745–5753. Checked.

[Zha96] K. Zhang, Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory, *J. Neurosci.* 16 (1996), 2112–2126.